An Interop Publication

CONNEXIONS®

The Interoperability Report

December 1991 Volume 5, No. 12

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

ISSN 0894–5926

From the Editor

It has been suggested that the Internet Transmission Control Protocol (TCP) will run at any speed over any media, present and future. Indeed, TCP's strength and historical success can largely be attributed to its hardware, software, and media independence. It is possible, however, to design a protocol that is specifically suited to a particular network environment. By "environment" is meant: the processor on which the protocol runs, the specific network media, and the types of end-system applications employed. Such considerations played an important part in the design of the *Xpress Transfer Protocol* (XTP). We asked one of XTP's principal architects, Greg Chesson, to give us an overview of this emerging protocol.

Daniel Dern gives us some impressions from INTEROP 91 Fall in the form of a "Dear Cliff" letter on page 9. This is followed by some snapshots from the show.

It is becoming clear that evolution to a ubiquitous worldwide Internet is not practical with the current IP addressing scheme, thus alternative solutions are being investigated. One possibility is to use IP's OSI equivalent *CLNP*. CLNP makes use of hierarchical Network Layer addresses, known as NSAPs. These addresses provide the flexibility needed to simultaneously solve two critical problems: (i) How to administer a worldwide address space; and (ii) How to assign addresses in a manner which makes routing feasible in a worldwide internet. Ross Callon, discusses how NSAP addresses could be used in the Internet.

At INTEROP 91 Fall, we were fortunate enough to have available a video teleconferencing system which connected San Jose, California to Geneva, Switzerland and the TELECOM 91 tradeshow. The system was put in place courtesy of US Sprint and Compression Labs, Inc. We used the link to "pipe in" Ellen Hancock's plenary address, and also for a session entitled "Who owns standards?" chaired by Carl Malamud. In Carl's session, it was announced that CCITT documents would be available, free of charge, over the Internet. Outraged by the high cost of standards documents, Carl has formed what he humorously refers to as the *Document Liberation Front*, and obtained permission from Dr. Pekka Tarjanne, Secretary General of the ITU, to convert and make available all 19,000 pages of the CCITT "Blue Book" series. Dr. Tarjanne was on hand (via the video link) to wholeheartedly endorse the project. In this issue, Carl describes his efforts in making CCITT (and hopefully soon ISO) documents available to anyone who wants them. We hope at least some of you will consider this an appropriate Christmas present from snowy Colorado. In any case, have good holiday, we'll be back with Volume 6 before you know it!

The Xpress Transfer Protocol (XTP)

by Greg Chesson,
Silicon Graphics, Inc. & Protocol Engines, Inc.

Introduction

Contemporary protocol research is influenced by developments in other technical fields. For example, *Very Large Scale Integration* (VLSI) techniques are commonly incorporated into the design of protocols, systems, and media—something that was not possible when many standard protocols were designed. Improved media in the 100Mbit/sec to 2000Mbit/sec range limit the usefulness of protocols designed for slower media. As a result, updated methods for flow, rate, and error control and new services receive attention. Applications such as advanced CAD and multimedia demand greater bandwidth and low latency interfaces. As a result, protocol mechanisms for providing bulk and media services, transactions, multicast, and network resource allocation become important.

Protocol evolution

Design of the *Xpress Transfer Protocol* (XTP) has been a 3 year process so far with roots traceable to work in the 1970's and 80's. XTP is a new protocol design that has been carried out by a multidisciplinary group including a VLSI design team, an operating system team, protocol architects, and real-time system experts. There has also been feedback from several software implementations. The evolution of XTP is defined by the interaction between these various interest groups over a period of years. The Final Report of *ANSI X3S3.3 Study Group on High Speed Network Protocols* [7] documents changing conditions and requirements that justify updating existing protocols or designing new ones.

XTP history

The chronology of XTP begins with work done at Bell Telephone Laboratories in the late 70's and early 80's by the author and others on lightweight protocols for the *Datakit*® network—a research project that can be described as the forerunner of today's ATM architecture.

During this same period several researchers at Bell Laboratories were practicing the art of rendering software algorithms into special-purpose hardware. First in this area was the chess machine, *Belle,* designed and built by Ken Thompson and Joe Condon. Similar methods were pioneered at Silicon Graphics in the 1980's with the development of the *Geometry Engine*® for accelerating graphics systems.

A common thread to these projects was that the developers did not simply apply VLSI techniques to existing algorithms: they had to transform their system algorithms to take advantage of and fit the constraints of hardware acceleration. The resulting systems are a hybrid of software, firmware, and special hardware. This suggests that VLSI techniques can be applied to any protocol, but will work best if the protocol can be adapted to the new environment. This is why hardware considerations have always played a part in XTP architecture.

Reaction to initial XTP designs fell into three categories: enthusiastic, we'll wait-and-see, and we'll stick-with-TCP (and so should you). We were obviously mandated to proceed.

In January 1988, a *Technical Advisory Board* (TAB) was established consisting of companies and research groups that would support the continued development of XTP funding, engineering reviews, testing and general advice. Although the TAB membership has changed during three years, it is active and strong and largely responsible for XTP achieving its present level of development.

Design goals

We wanted a protocol that would facilitate the design of supporting VLSI and which would be suitable for operation at gigabit speeds. A simple design based on fixed-size headers and, trailers and no variable-size options was indicated. On the other hand we didn't want a simplistic protocol that might not be able to provide the services needed by modern distributed systems.

It was decided that several service models—message, stream, transaction, bulk, controlled rate, and multicast—were needed. It was observed that with the exception of multicast, these services were all the same: deliver some bits between point A and point B. The difference between these services is that they represent different distributions of data traffic and frequency. Consequently we could try to design a protocol that could supply these different services, including multicast, with a single suite of mechanisms.

A modern protocol should anticipate an environment populated by both switching and non-switching media. The protocol should operate well with either kind of network and also provide a means for coupling between dissimilar media. In addition it was decided that a new protocol should eliminate congestion at both end nodes and interior network nodes by incorporating active control policies. These should include active protocol exchanges between internal nodes and end nodes to establish rate control and resource allocation.

An early concept in XTP design was, and still is, the notion of "real-time" processing. The idea is to pass data units, or packets, through a pipeline of processing engines. The bidirectional pipeline would connect to a network at one end and to a host computer system, router, or other device at the other end. The data rate through the pipeline should equal or exceed the media data rate. If protocol and system I/O processing can be completed within the pipeline, at the media data rate, then the network and I/O processing can be termed "real-time."

Address translation

Address translation, whereby address bytes from a header are used as the lookup key to a database lookup procedure, can be the most time consuming component of input protocol processing.

Address lookups can be accelerated using caching schemes. Although these are easier to implement in software than hardware, techniques such as Van Jacobson's header prediction scheme for TCP can be helpful. Unfortunately, a header prediction technique does not reduce software processing overhead when the prediction fails. Applications like The X Window System cause a server to receive random short messages from different systems. If the header cache size is increased to accommodate the most active connections, then the cache lookup problem has about the same overhead as the original address translation problem.

For these reasons XTP evolved to the KEY-based scheme employed today where there are three items extracted from a packet for address translation: MAC source address, 32-bit KEY field, and 32-bit ROUTE field. Once an initial handshake has been completed between end systems, address translation reverts to a single index-based table lookup. This has about the same overhead as a header prediction scheme, but works on every packet.

The initial packet that sets up an XTP connection, called a FIRST packet, contains a variable-length *address segment*. The syntax of an address is specified by a *type* field.

continued on next page

The Xpress Transfer Protocol *(continued)*

XTP specifies several address formats for compatibility with the Internet world, the ISO universe, and other addressing plans. This notion of a *parameterized address* may be unique to XTP—it enables the addition to existing networks without introducing an additional address plan.

The design of Protocol Engine hardware pipelines progressed from trie-based schemes to hash-based schemes. Hardware was designed to extract address bytes from a passing packet according to protocol type, and generate a CRC-based hash function. A specialized processor is used to complete the hash lookup. The processor has special instructions for reducing the time required for comparing binary strings, and a FIFO ensures that the occasional long hash chain does not cause input processing to drop packets.

Note that the hash lookup is needed only for FIRST packets or until a so-called KEY exchange has been completed. Non-FIRST packets are processed with a table lookup. The result is that hashing is needed for only a small percentage of packets.

Framing

XTP incorporates fixed-size headers and trailers aligned on multiples of eight bytes. Eight byte alignment is provided for programming efficiency when the protocol is implemented in software on a 64-bit machine. It also anticipates future Protocol Engine designs that might have internal 64-bit data paths.

The headers and trailers do not change size when options are selected. Rather, the bitfields for all protocol mechanisms are present at all times.

A field in the XTP header is defined as a *length* field. This makes packet processing simpler to implement in VLSI, and eliminates reverse parsing logic from hardware pipelines. The length field and related simplifications are explained in Addendum 1a to Revision 3.5 of XTP.

Internetworking

XTP defines a cut-through switching technique for the routing of packets between networks. The technique is similar to setting up a virtual circuit through a network of switches. The difference is that each XTP "switch" is allowed to participate in the end-to-end rate-based flow control. Thus each intermediate node, or "switch," is able to exert its own resource allocation policy on the end nodes that are communicating via the intermediate node. This technique can eliminate congestion at both intermediate and end nodes. It also seems to be an important prerequisite for establishing isochronous flows through an internet.

SORT field

XTP has incorporated a priority mechanism for both input and output ordering. Prior to Revision 3.5 the field could be interpreted in two ways: either as a 32-bit static priority, or as a 32-bit time deadline. More hours of discussion went into SORT field semantics and related real-time issues than almost any other aspect of XTP design.

Unfortunately no plausible method for converting between the two representations, or specifying how the two mechanisms could interoperate, was discovered by those interested in this aspect of XTP. For these reasons the SORT field definition was simplified to being just static priority. The issue of representing more advanced scheduling metrics in the protocol remains a research issue. It is hoped that we will be able to return to this issue after gaining experience with XTP in advanced real-time systems.

Record structure

Protocols are often used to transmit sequences of records or messages whose binary image is then reconstructed at a receiver. Protocols are also used to provide a byte stream service. What we wanted in XTP was a byte stream service that would also provide message facilities adequate for record marking, encapsulation, and convergence protocols.

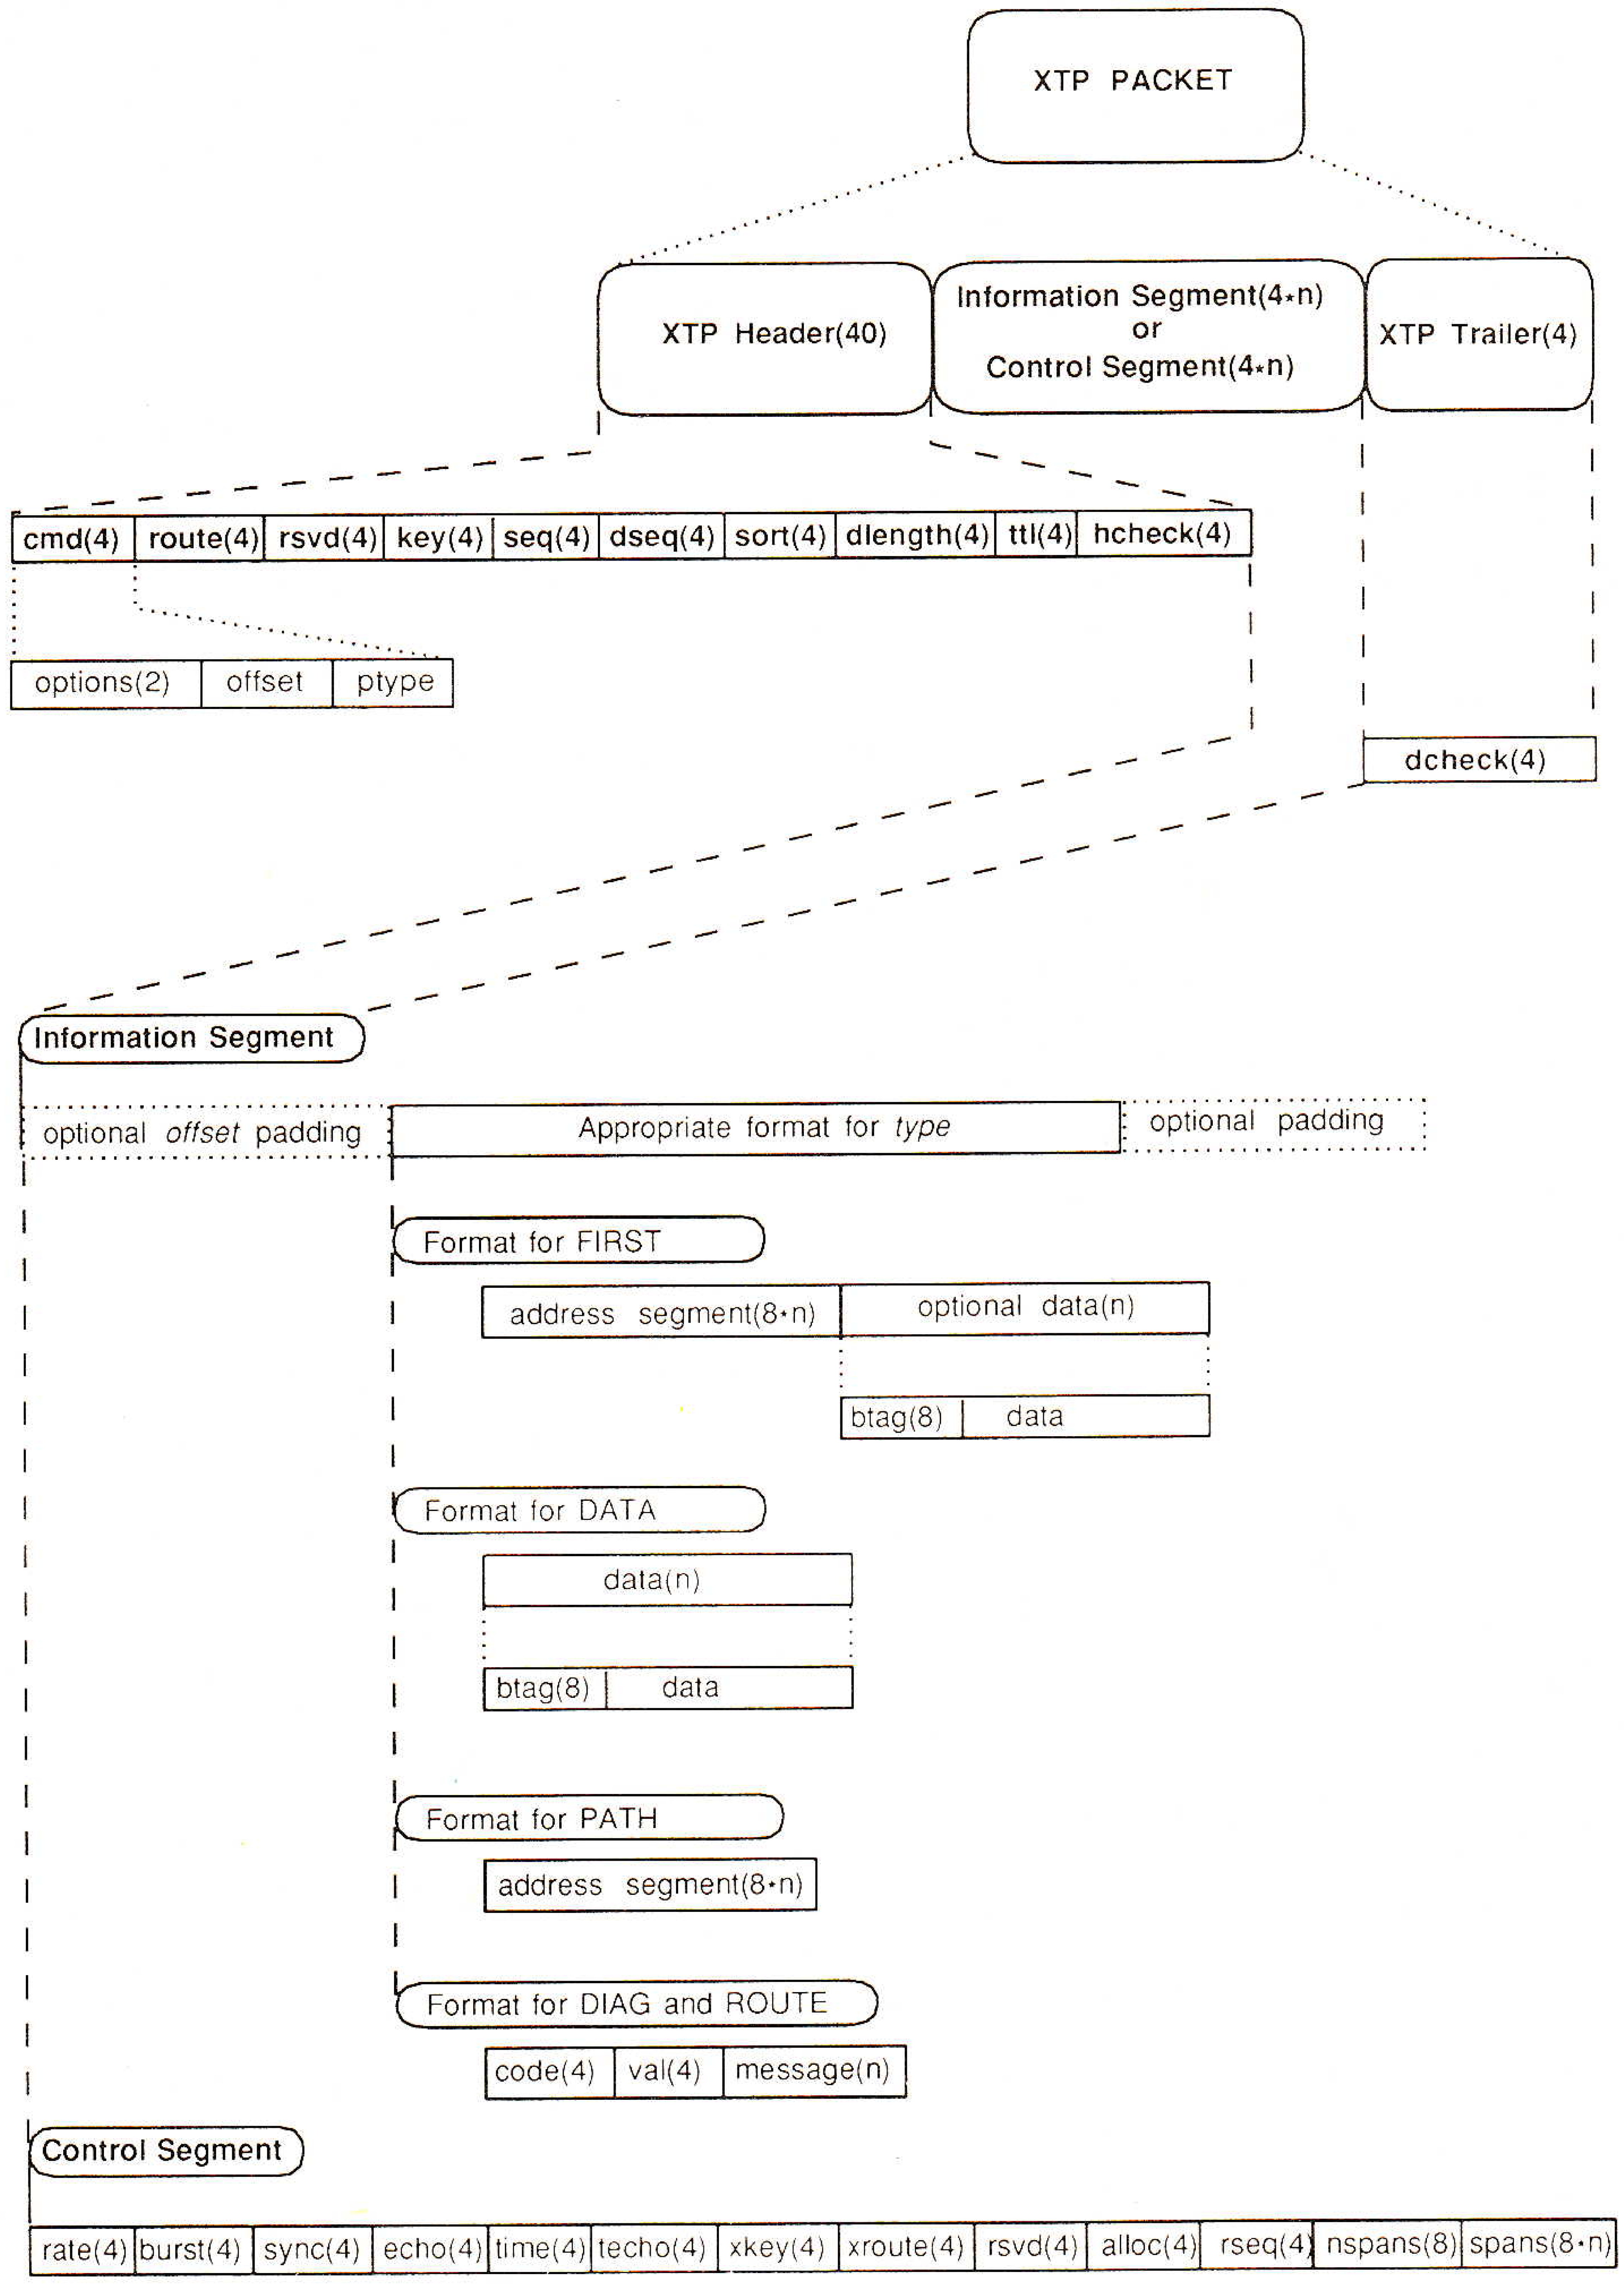

Figure 1: XTP Packet Format

We define *record marking* as the facility to send a sequence of arbitrary byte length messages on a byte stream protocol and to reconstruct the original messages at the receiving end. We define *encapsulation* as the ability to transport protocol data units produced by a protocol such as TCP/IP or ISO TP4 as data units in XTP. The transported data units would be carried through an XTP subnetwork as data and injected into a native TCP/IP network or ISO TP4 network at an XTP gateway.

continued on next page

The Xpress Transfer Protocol *(continued)*

We defined a *convergence protocol* as one that provides encapsulation as just described, but also provides the service of the encapsulated protocol. An example of convergence is provided by a modified NFS implementation, which normally used UDP, but instead multiplexes NFS transactions to a server over a shared XTP association. The method for representing the NFS transactions and delivering them to the NFS software is an example of a convergence protocol.

Message service can be provided with a stream protocol by two basic methods: either impose a record structure on the stream, or use a record delineating mechanism provided by the underlying transport. The imposition of a record structure on the stream is easy, but potentially costly since it requires scanning and parsing the data stream. Reliable stream protocols seldom expose the underlying frame structure to the stream user. One notable exception is the *Sequenced Packet Protocol* (SPP) defined as part of the XNS protocol suite. This protocol provides an 8-bit Data Stream Type that is setable on output and delivered on input.

XTP provides a method for marking the data stream. The EOM bit in the trailer, borrowed from the XNS mechanism of the same name, accomplishes this. For encapsulation the protocol provides for a 64-bit *tag* field at the beginning of the application data segment. The presence of the field is indicated by the BTAG bit in the header.

Checksums

The XTP checksum was always intended for hardware implementation in a parallel data path. This meant that checksum logic could easily become part of a host DMA engine. It also meant that a software implementation could do 32-bit or 64-bit fetches from memory instead of fetching one or two bytes per algorithm step.

Early versions of the XTP checksum were parallel, but not portable. That is, a little-endian CPU and a big-endian CPU could compute the checksum quickly, but would get different results. The choices for correcting the problem were to force one of the CPU types to reorder bytes before checksumming, or to change the algorithm. The algorithm was therefore changed to its present form, called CXOR.

Sequence Numbers

Most quantities in XTP control structures are 32-bit values. In older versions a provision was made for also having 64-bit values.

The fixed-size structures of XTP allow for a stateless conversion between 32-bit and 64-bit implementations. It was shown that 32-bit implementations could interoperate with 64-bit versions through a translating gateway or by a simple conversion process.

The trick is that the fixed size headers and alignment rules allow a simple doubling or halving of header and trailer sizes. An additional bit is needed in the header to direct a 64-bit implementation to truncate its arithmetic to 32-bits. This is one example of a mechanism that was removed from the protocol, but which could be revived when multi-gigabit networks are deployed.

Error control

Experience with Datakit where the hardware provides a switched order-preserving packet stream suggested that a *go-back-N* retransmission model would be sufficient for high speed applications. This simple model is not best for all applications.

The pathological cases occur over slow media, in environments that can reorder packets, and in high bandwidth-delay environments. In these cases a selective retransmission scheme is useful whereby only error frames are retransmitted.

The first definitions of XTP used only go-back-N. When a selective retransmission method was added, the model incorporated the notion of sequence number *spans*. Each span is an ordered pair of sequence numbers. Spans originally designated missing data, or sequence *gaps*, detected by a receiver. Later the spans were shifted so that instead of representing gaps of missing data, they represented the islands of correctly received data.

Care was taken throughout to guarantee that an implementation would provide go-back-N service and still interoperate with another implementation that provided selective retransmission based on spans. Since both the algorithm and the number of spans can be chosen by the implementor, this part of the design has been called *selectable retransmission.*

Multicast

Implementation groups began to experiment with multicast in an effort to characterize and improve the performance of a multicast stream. This work led to an implementation technique for flow control called the *bucket algorithm*. This procedure and other multicast techniques are documented in Revision 3.5. What has not been documented before is the design choices made in non-multicast that enable the multicast mode.

First there is a MULTI bit in the header that is set on multicast output and tested on input. This bit simplifies special case testing where multicast processing is different from unicast. More important are the *sync* field and the SREQ bit. The SREQ bit requests an immediate state acknowledgement from the receiver. Not all protocols have this mechanism. XTP uses it for many things: determining round trip times, determining when to retransmit, and for generating a flow of state acknowledgements that are controlled by the protocol sender. The ability to control acknowledgements lets XTP divide the output flow into time slices called *buckets*. The sync field is used to identify each bucket and to identify any acknowledgements that were caused by an SREQ sent during a bucket interval.

This kind of information is overkill for managing the flow of a single stream, but in XTP multicast the flow control processing must be spread out over time, i.e., some number of buckets, to get adequate error control. The ability to manage the flow of control messages and identify output epochs are critical components of the multicast procedure. By designing them into the unicast procedure we avoid having to add redundant functions to the base protocol for supporting multicast.

Acknowledgements

The architecture has been influenced by many other designs including Delta-T, VMTP, Datakit protocols, Netblt, Blast protocols, ATM, TCP/IP, GAM-T-103 [1], the Cambridge Ring, and XNS. More people have contributed to the evolution of XTP than can be named here. The greatest concentration of effort has been though the TAB group and its Research Affiliates, the ANSI X3S3.3 High Speed Protocols Study Group, and the design and implementation teams at Silicon Graphics Computer Systems and Protocol Engines, Incorporated. Support and encouragement to pursue this work is gratefully acknowledged.

continued on next page

The Xpress Transfer Protocol *(continued)*

References

[1] Ministre de la Défense, République Française, "Military Real Time Local Area Network," GAM-T-103, February 9, 1987.

[2] Protocol Engine and XTP are registered trademarks of Protocol Engines, Incorporated.

[3] G. Chesson, "Protocol Engine Design," in USENIX Conference Proceedings June 1987, p. 209, also "The Protocol Engine Project" in *UNIX Review,* September, 1987, also P-Engine Protocol Definition, September 1987, Protocol Engines Inc., 1900 State Street, Santa Barbara, CA 93101.

[4] "XTP Definition, Revision 3.5 plus Addendum 1a," September 10, 1990, Protocol Engines Inc., 1900 State Street, Santa Barbara, CA 93101.

[5] William D. Schwaderer, "XTP in VLSI—Protocol Decomposition for ASIC," Proceedings of the 15th IEEE Conference on Local Computer Networks, Minneapolis, Minnesota, September 1990.

[6] W. Doeringer, et. al., "A Survey of Lightweight Transport Protocols," IBM Research Report, RZ 1980 (number 70152), May 17, 1990.

[7] James Moulton, "Final Report for Project 753, ANSI document X3S3.3/90-201R, Results of the Study on Very High Speed Network Protocols," December 11, 1990.

[8] Ian Wakeman, "Emerging High Speed Transfer Protocols," IEEE Telecom Conference Proceedings, March 1991.

[9] Harry Rudin and Robin Williamson, editors, *Protocols for High-Speed Networks,* North-Holland, 1989.

[10] G. Chesson, "XTP/PE Design Considerations," in *Protocols for High-Speed Networks,* [9].

[11] Robert Sanders, "XTP Tutorial," University of Virginia Computer Science Department.

[12] T. Mason, "VMTP: A High Performance Transport Protocol," *ConneXions,* Volume 4, No. 6, June 1990.

[13] S. Deering, "IP Multicasting," *ConneXions,* Volume 5, No. 2, February 1991.

[14] P. McKenney, "Congestion Avoidance," *ConneXions,* Volume 5, No. 3, March 1991.

GREG CHESSON is currently Chief Scientist at Silicon Graphics Inc., and Founder and Chief Technical Officer of Protocol Engines Inc. Chesson holds a Ph.D. degree in Computer Science from the University of Illinois, 1977. He was a member of the Technical Staff at Bell Laboratories, Murray Hill, from 1977 until 1982. At Bell Labs he worked in Laboratory 127 which released Editions 5, 6, 7, and 9 of the UNIX operating system to the commercial world. His contributions to the Seventh Edition of UNIX include the packet driver protocol in *uucp,* the UNIX MPX multiplexor, and numerous UNIX internals. Chesson designed and implemented the Datakit protocol suite, now in common usage at BOC central exchanges. He designed and implemented the B-machine which was a TTL prototype for a protocol engine on Datakit. The B-machine is considered to be the first example of a transport protocol executable directly in hardware. At Silicon Graphics Inc. since 1982, Chesson implemented XNS protocols for early SGI workstations and for UNIX and VMS operation systems. As Chief Scientist, he has contributed to numerous UNIX, network, and graphics architecture projects. In 1987, Chesson founded Protocol Engines Inc. and began developing the Protocol Engine® design effort now in progress.

"Dear Cliff": A Report From INTEROP 91 Fall

by Daniel P. Dern

From: Daniel Dern <ddern@world.std.com>
To: Cliff Stoll <astro@cfa253.harvard.edu>
Subject: What You Missed At INTEROP 91 Fall
Date: Halloween, 1991
Cc: Ole Jacobsen <ole@interop.com>

Dear Cliff,

Since you had to miss this year's INTEROP show (OK, having to go to Japan on a book-promotion tour for *The Cuckoo's Egg* is a good excuse :–), I thought I'd fill you in on some of what you missed. I know you can read about most of the big announcements and other hoo-hah in the trade press, so here's a round-up of events and stuff that might interest you wearing your various hats as scientific end-user, Internet member, security maven, and fun-loving guy.

35,000

First, the numbers. Headcount for INTEROP 91 Fall was about 35,000—50% more than last year's, about four times 1989's. The move from San Jose to San Francisco's Moscone Center for 1992 comes none too soon. I'm glad I made my hotel reservations months in advance. I know that INTEROP filled up the Convention and Civic Centers, plus used the main hall in the Center for Performing Arts, and most of the Fairmont Hotel. Over 250 vendors exhibited, and attendees had to steer a course through these, the on-going tutorials, sessions BOFs, and evening parties.

This Interoperability biz has gotten even more serious than last year, if dress code is any indicator. More "suits"—male and female. Fewer blue jeans and t-shirts (although still a clear showing), and I didn't spot anyone in bare feet, despite the warm weather.

Shownet

Once again, the Shownet squad deployed the INTEROP show floor internetwork for connecting all the show floor exhibits: over thirty miles of 10Base-T unshielded twisted pair and three-pair FDDI "orange garden hose" comprising four backbones and twenty five "ribs" across two buildings, running TCP/IP and using OSPF for the primary routing protocol (and static routes from the booths to the routers), and supporting about 350 subnets and 4,000 devices. The Shownet gang rolled it out mostly in an eight-hour marathon starting late Friday evening.

Alexander Latzko, a telecom analyst from Rutgers University, ballparks Shownet as the equivalent of what goes into a twenty-story high-tech skyscraper.

Karen Auerbach (Epilogue Technology) noted, "The Shownet went up much easier this year than last, and stayed up more." She credits much of this success to the 'wiring party' from July where "we lay the network wiring and then wrap it in spools and store it in a warehouse until October, giving Shownet the distinction of being the only network that gets stored in mothballs." (And perhaps the first network to be mothballed *before* being used :–)

SNMP

"The demonstrations for SNMP this year represented much more what a network was like," Auerbach also notes. "We got vendors to agree to put 'naughty' equipment, misconfigured or otherwise giving problem behaviors, which gave people a much more realistic view of what network mangers see and how SNMP helps to diagnose and isolate the problems.

continued on next page

A Report From INTEROP 91 Fall *(continued)*

We had 44 participating exhibitors—including many for whom this was their first INTEROP and their first SNMP product. This is exciting, because many of these companies, like BICC, Du Pont and Sumitomo are well known names in the networking world, but not necessarily associated with the TCP/IP market or LAN arena."

Oracle, by its own admission, had "the hottest exhibit on the show floor." Literally—their exhibit accounted for one third of the total electrical consumption. As part of their demo showing connectivity across different vendor and network platforms, they cleaned out their basement, and brought about fifteen tons of computers, everything from a ten-year-old Xerox Dandelion to an Ncube, Sequent and Auspex, taking up 600 square feet of raised computer floor, and a 100 amp feed at 480 volts. Jack, have that yard sale!

Applications

The gangs from Epilogue, FTP and TGV who last year brought us the Internet Toaster and CD Player were at it again, showing real-world applications of Internet technology including:

- An SNMP-controlled Lego construction to load and retrieve toast.
- A PC-controlled Lego plotter, (see photo on left) run by a VAX doing RPCs to a network controller under Beame & Whiteside TCP with SunOS RPCs, controlled with TGV's *XView* interface, demonstrating RPC-based client-server connectivity and how inexpensive the platforms can be—cheap enough to dedicate a PC card in Santa Cruz to monitor surfing weather.
- Simon Hackett, back from Australia, with the voice-packetizing *Etherphone,* stuffing audio signals into UDP packets, connecting phone sets across the floor, and letting us listen to a radio tuner in Melbourne. Kind of like ISDN in reverse, I guess. A great way to make use of that extra IP bandwidth.
- SNMP leader Jeff "That Dawg [Dog] Can Hunt" Case and SNMP Research brought encryption-based Secure SNMP, personified by the "One-Man Dog."
- Geoff Goodfellow brought along *RadioMail*, a service that lets you pick up Internet and other e-mail on radio pagers, and send as well as receive from radio-modem equipped portable computers. (If this had been available a few years sooner, the *Cuckoo's Egg* might have been caught one chapter sooner—maybe.)
- Plus there was lots of PC and Mac stuff—TCP/IP under Microsoft Windows, Mac servers accessed from X Windows—and FDDI over *copper* at 100Mbps...endless neat stuff!

The Great IGP Debate

One of the highlights (for many of us, anyway) was "The Great Interior Gateway Protocol Debate," where OSPF, IS–IS, Ships in the Night and Integrated Routing squared off in a semi-debate. Chaired with deadly wit by Lyman Chapin, Milo Medin, Radia Perlman, Ross Callon and Dave Clark (substituting for Noel Chiappa) held somewhere between 800 to 1000 people enthralled in a no-holds-barred, hour-plus session, after which Ole committed INTEROP to "at least one non-serious event per conference."

While many, perhaps most, of the attendees came to attend tutorials and sessions and to stroll the floor, there was a strong presence and sub-text of the Internet (which gave us TCP/IP, OSPF, SNMP and lots more of these now-salable technologies) itself.

Luminaries

Internet luminaries, developers and other notables could be seen strolling around: Vint Cerf, Einar Stefferud, Susan Estrada, Mike Padlipsky, Steve Kent, Marshall Rose, Jon Postel, Yakov Rekhter, Jack Haverty, Erik Fair, Bob Braden...the usual gang of suspects. Mitch Kapor was there on behalf of the Electronic Freedom Foundation, and commented, "It feels like a baby Comdex." (That overly big computer show they hold in Los Vegas every year.)

The Internet itself was there, not only through the rows of color X terminals letting us telnet through the Shownet server to read our e-mail, but also via booth exhibits by IP regional and commercial providers including ANS, BARRnet, CERFnet (sporting "Truth, Justice and the Internet Way" comics and buttons), PSI, Uunet, plus MERIT and the NISC (formerly SRI NIC). (You'd think there would also be some generic "Internet" booth/display. Maybe next year.) Joyce Reynolds from SRI ran sessions on INTEROP User Services; There was a late-night BOF for the IP commercial network providers. The *Archie* archive server developers from McGill were there...and something that should interest you wearing your astronomer's hat: there's a WAIS (*Wide Area Information Server*) with graphic star images you should be able to telnet to under X windows, somewhere in the Carolinas. I tried to access it from e-mail alley, but ran out of time. (Msg me and I'll send you the hostname.)

Changing audience

John "Toaster" Romkey sees the new INTEROP attendees as having more interest in using the technologies and products than being the developers and pioneers that comprised the first few INTEROPs. "It's now the 'gang' plus second cousins plus distant relatives and strangers...For example, in the packet driver BOF, we were hoping to iron out some details of the coming implementation. But only about 10% of the folks who showed up talked. The rest were there to learn."

One final observation: having now attended three year's worth of INTEROPs, I have seen technologies evolve—within two years—from announcement to first products to well-rolled-out commercial offerings. Examples: OSPF, PPP and SNMP, announced at INTEROP 89, demoed in 90—running and for sale in 91. And we know it works, because we see it working.

Sorry you had to miss INTEROP this year—I know we missed having you there.

—*Daniel*

A frequent contributor to *ConneXions,* **DANIEL P. DERN** is a Watertown, Mass.-based writer specializing in technology, science and business. He also writes computer humor, science fiction, and musical theater.

1991 INTEROP Achievement Award Winners

The *INTEROP Achievement Award* is presented to those customer organizations that make the most effective use of internetworking technology to further their own specific business aims. The award is given to one organization from each of four business categories. This year's winners were presented on October 9th at a special ceremony. The winners are:

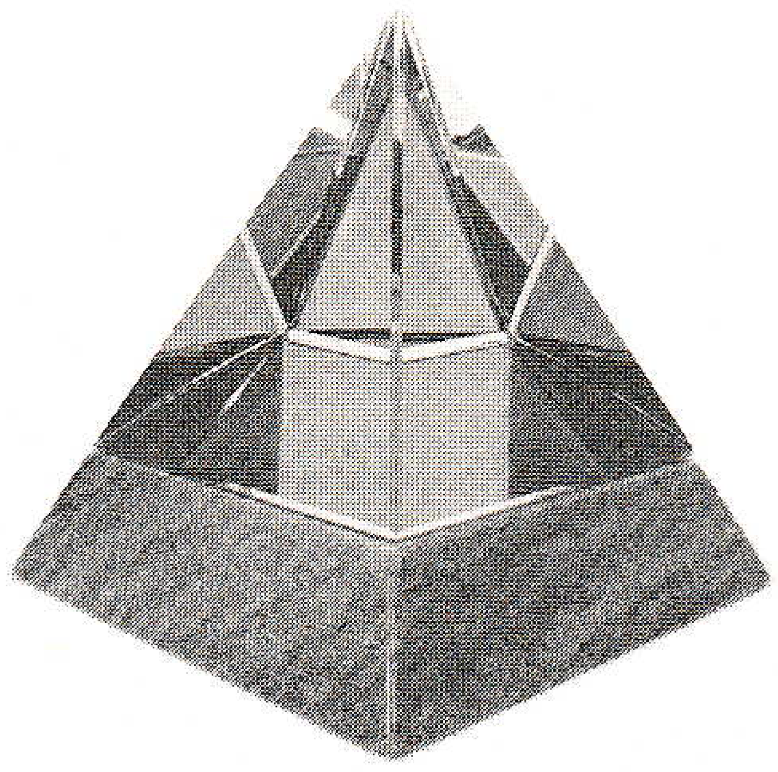

- *Service:* Fidelity Investments
- *Education:* State University of New York
- *Manufacturing:* The Foxboro Company
- *Government:* Sandia National Laboratories


INTEROP
INTEROP

Snap-shots

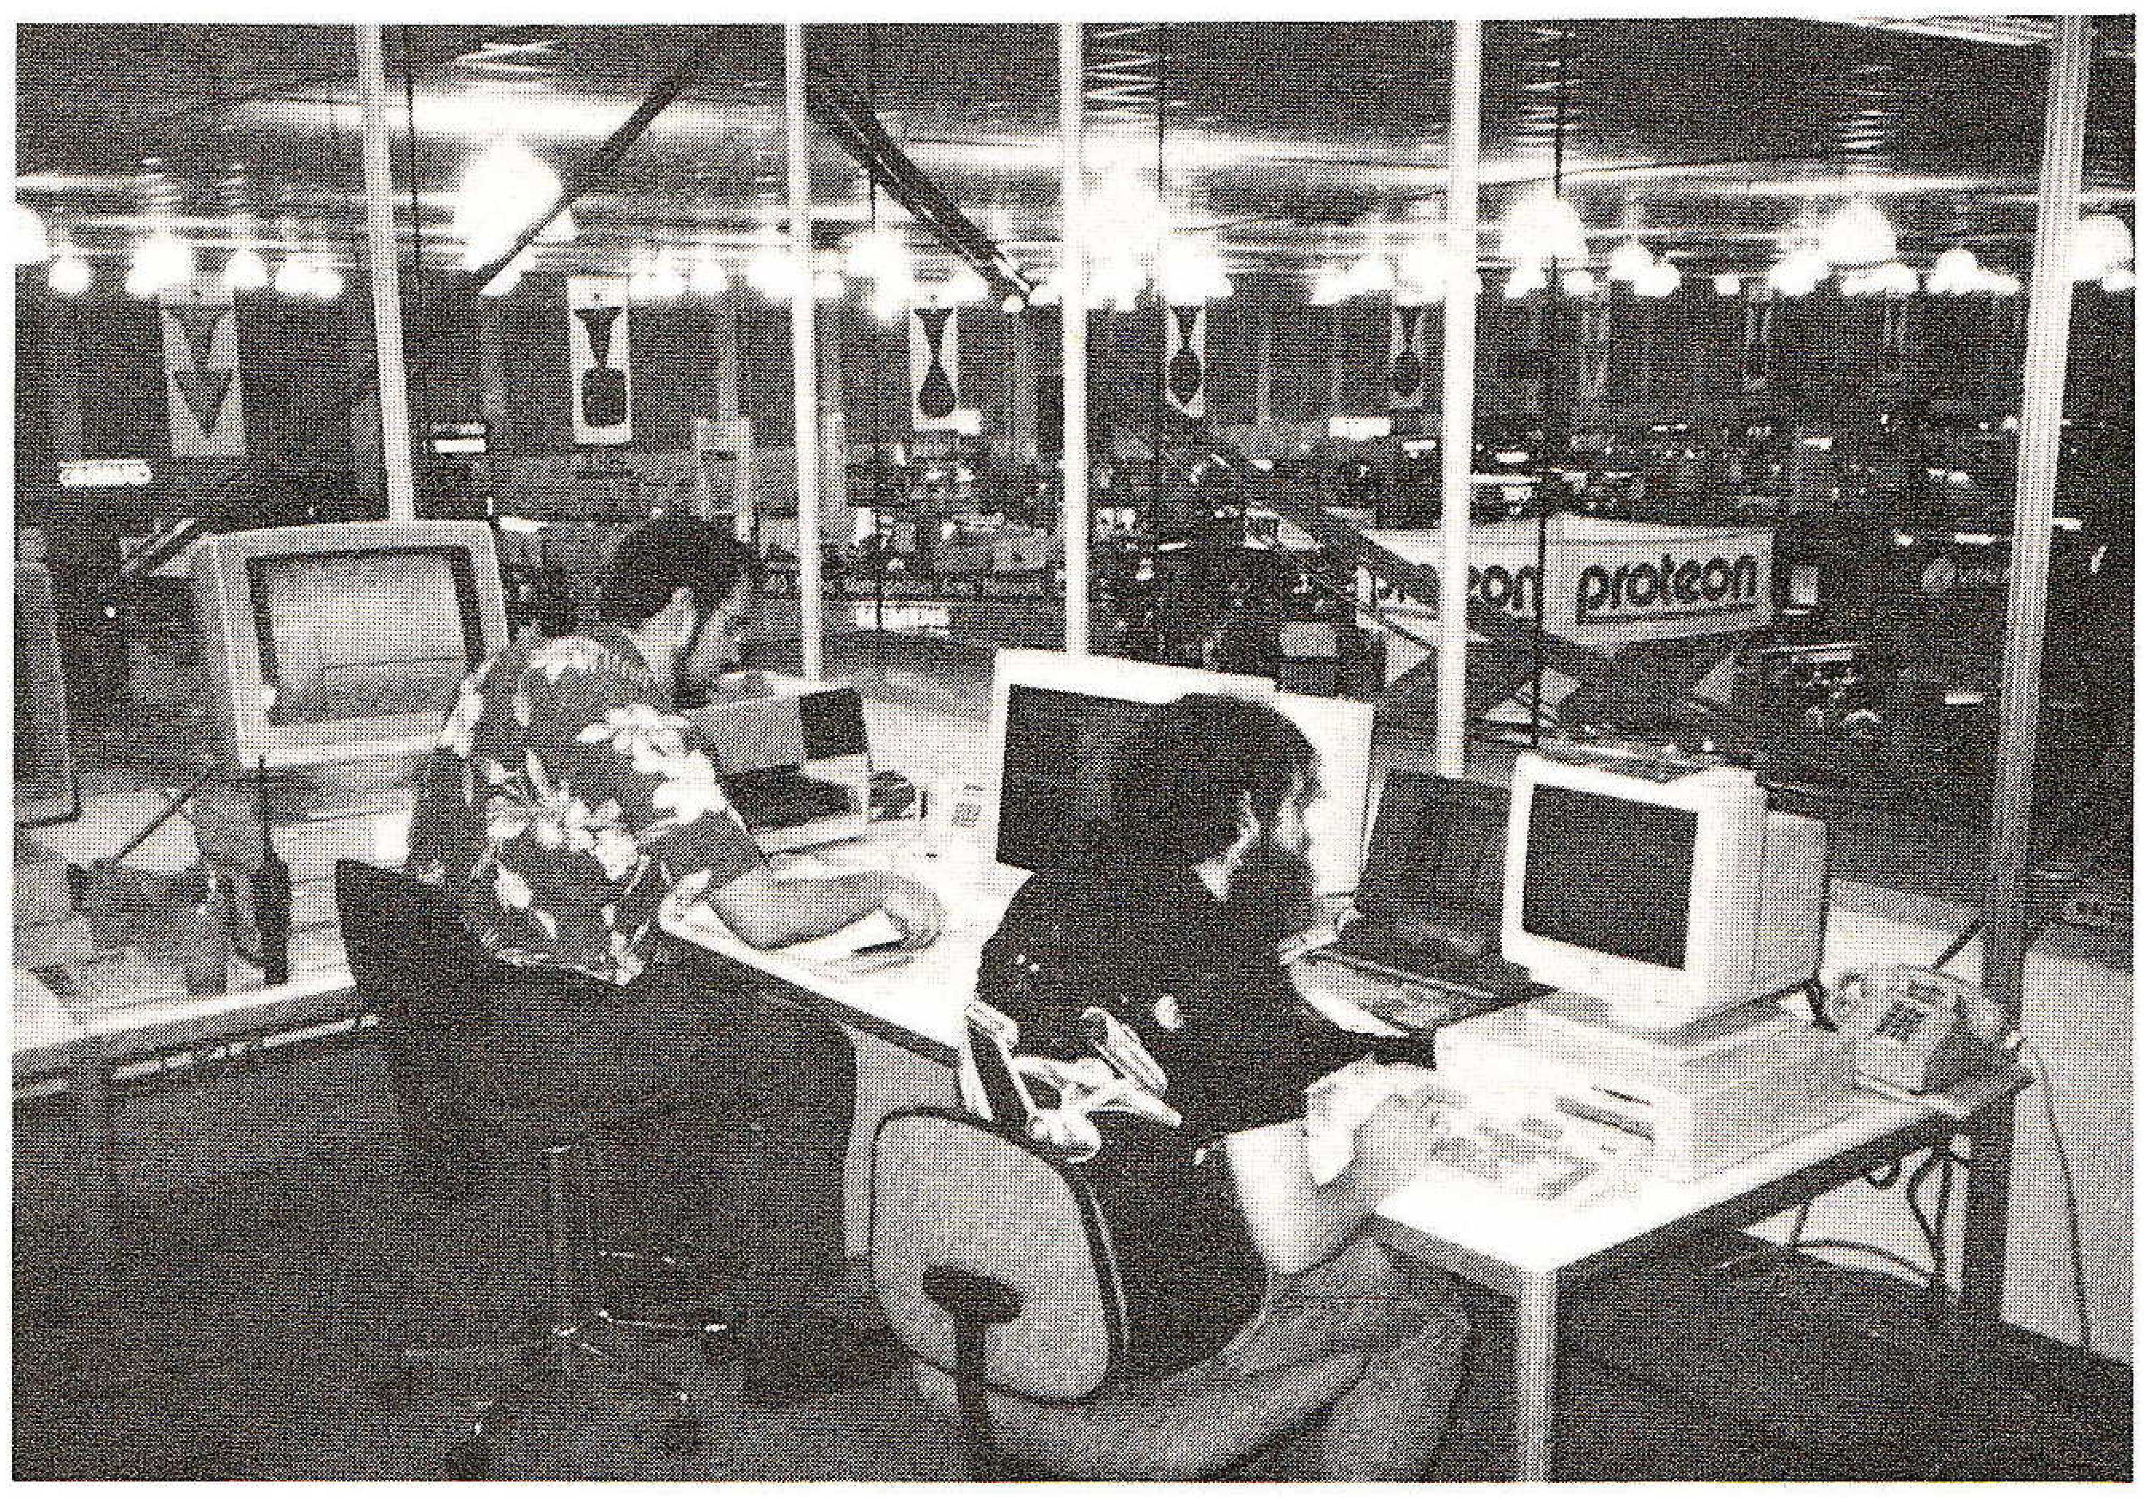
proteon


INTEROP
Visit the
Solutions Showcase
Demonstrations
Participants
Look for these Logos on the Exhibit Floor
SMDS
Token Ring
FRAME RELAY
OSI
ISDN


3Com
AUSPEX
1200
ORACLE
1300
1400
CMC

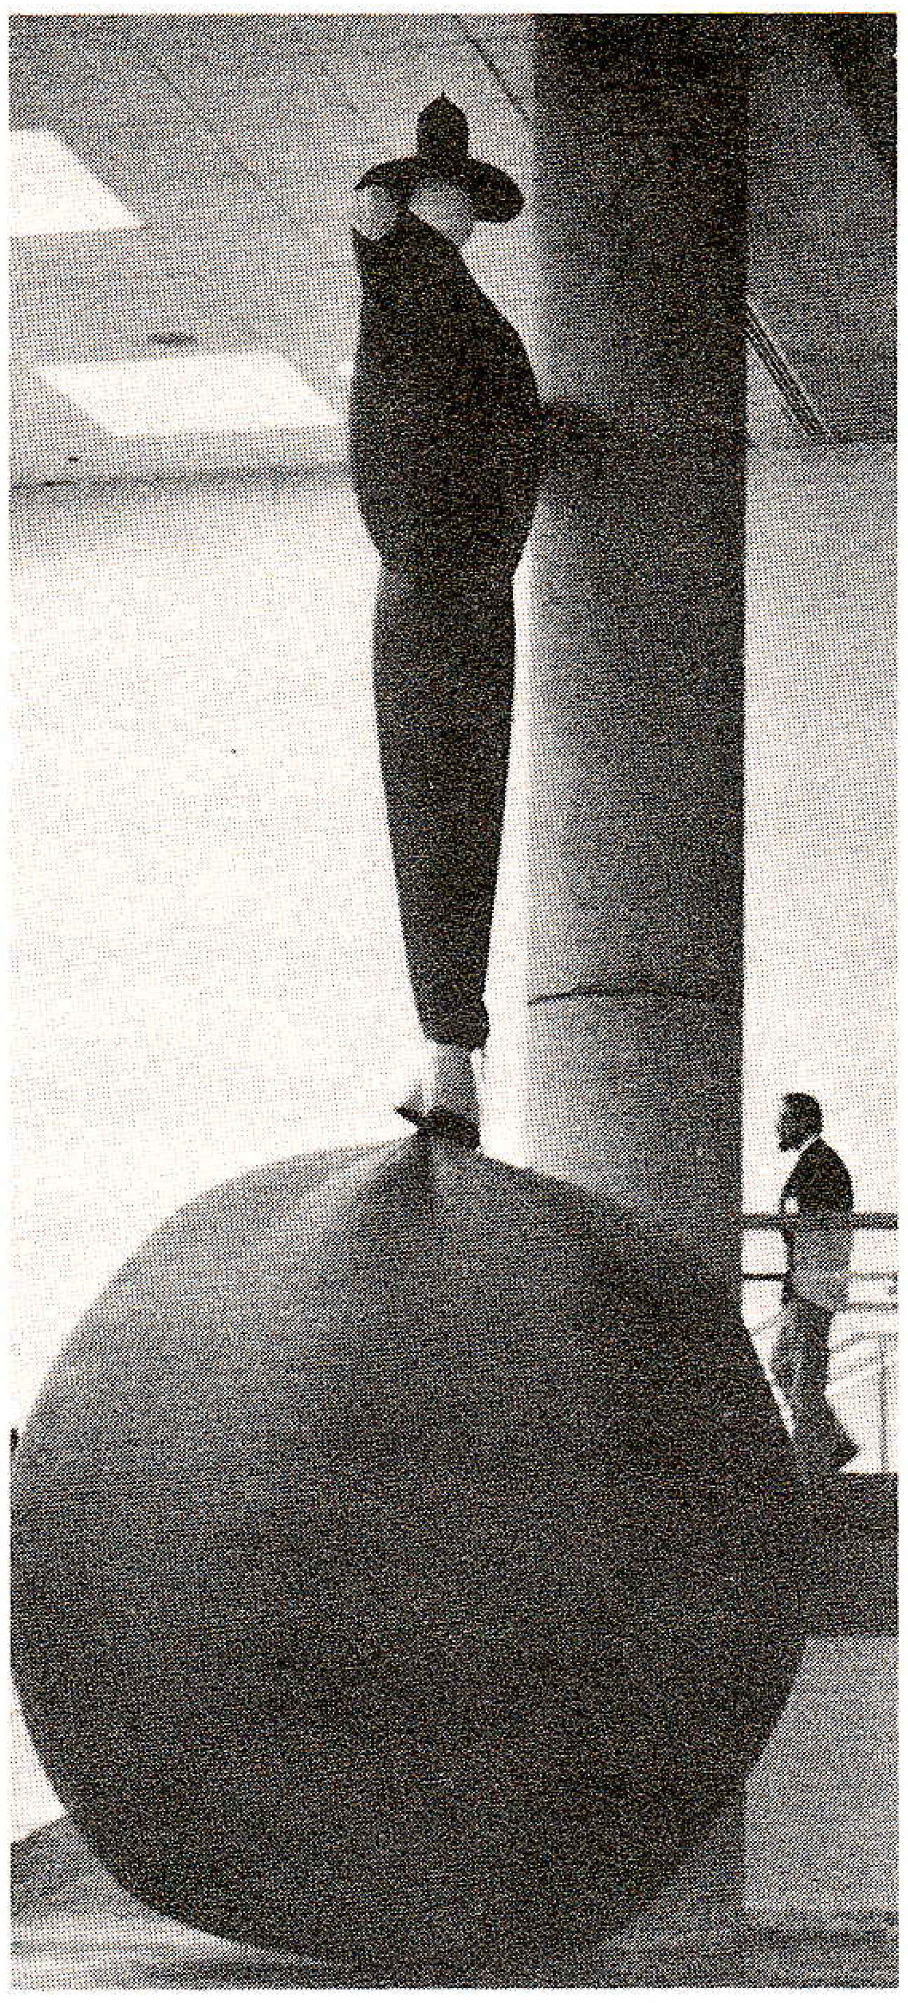

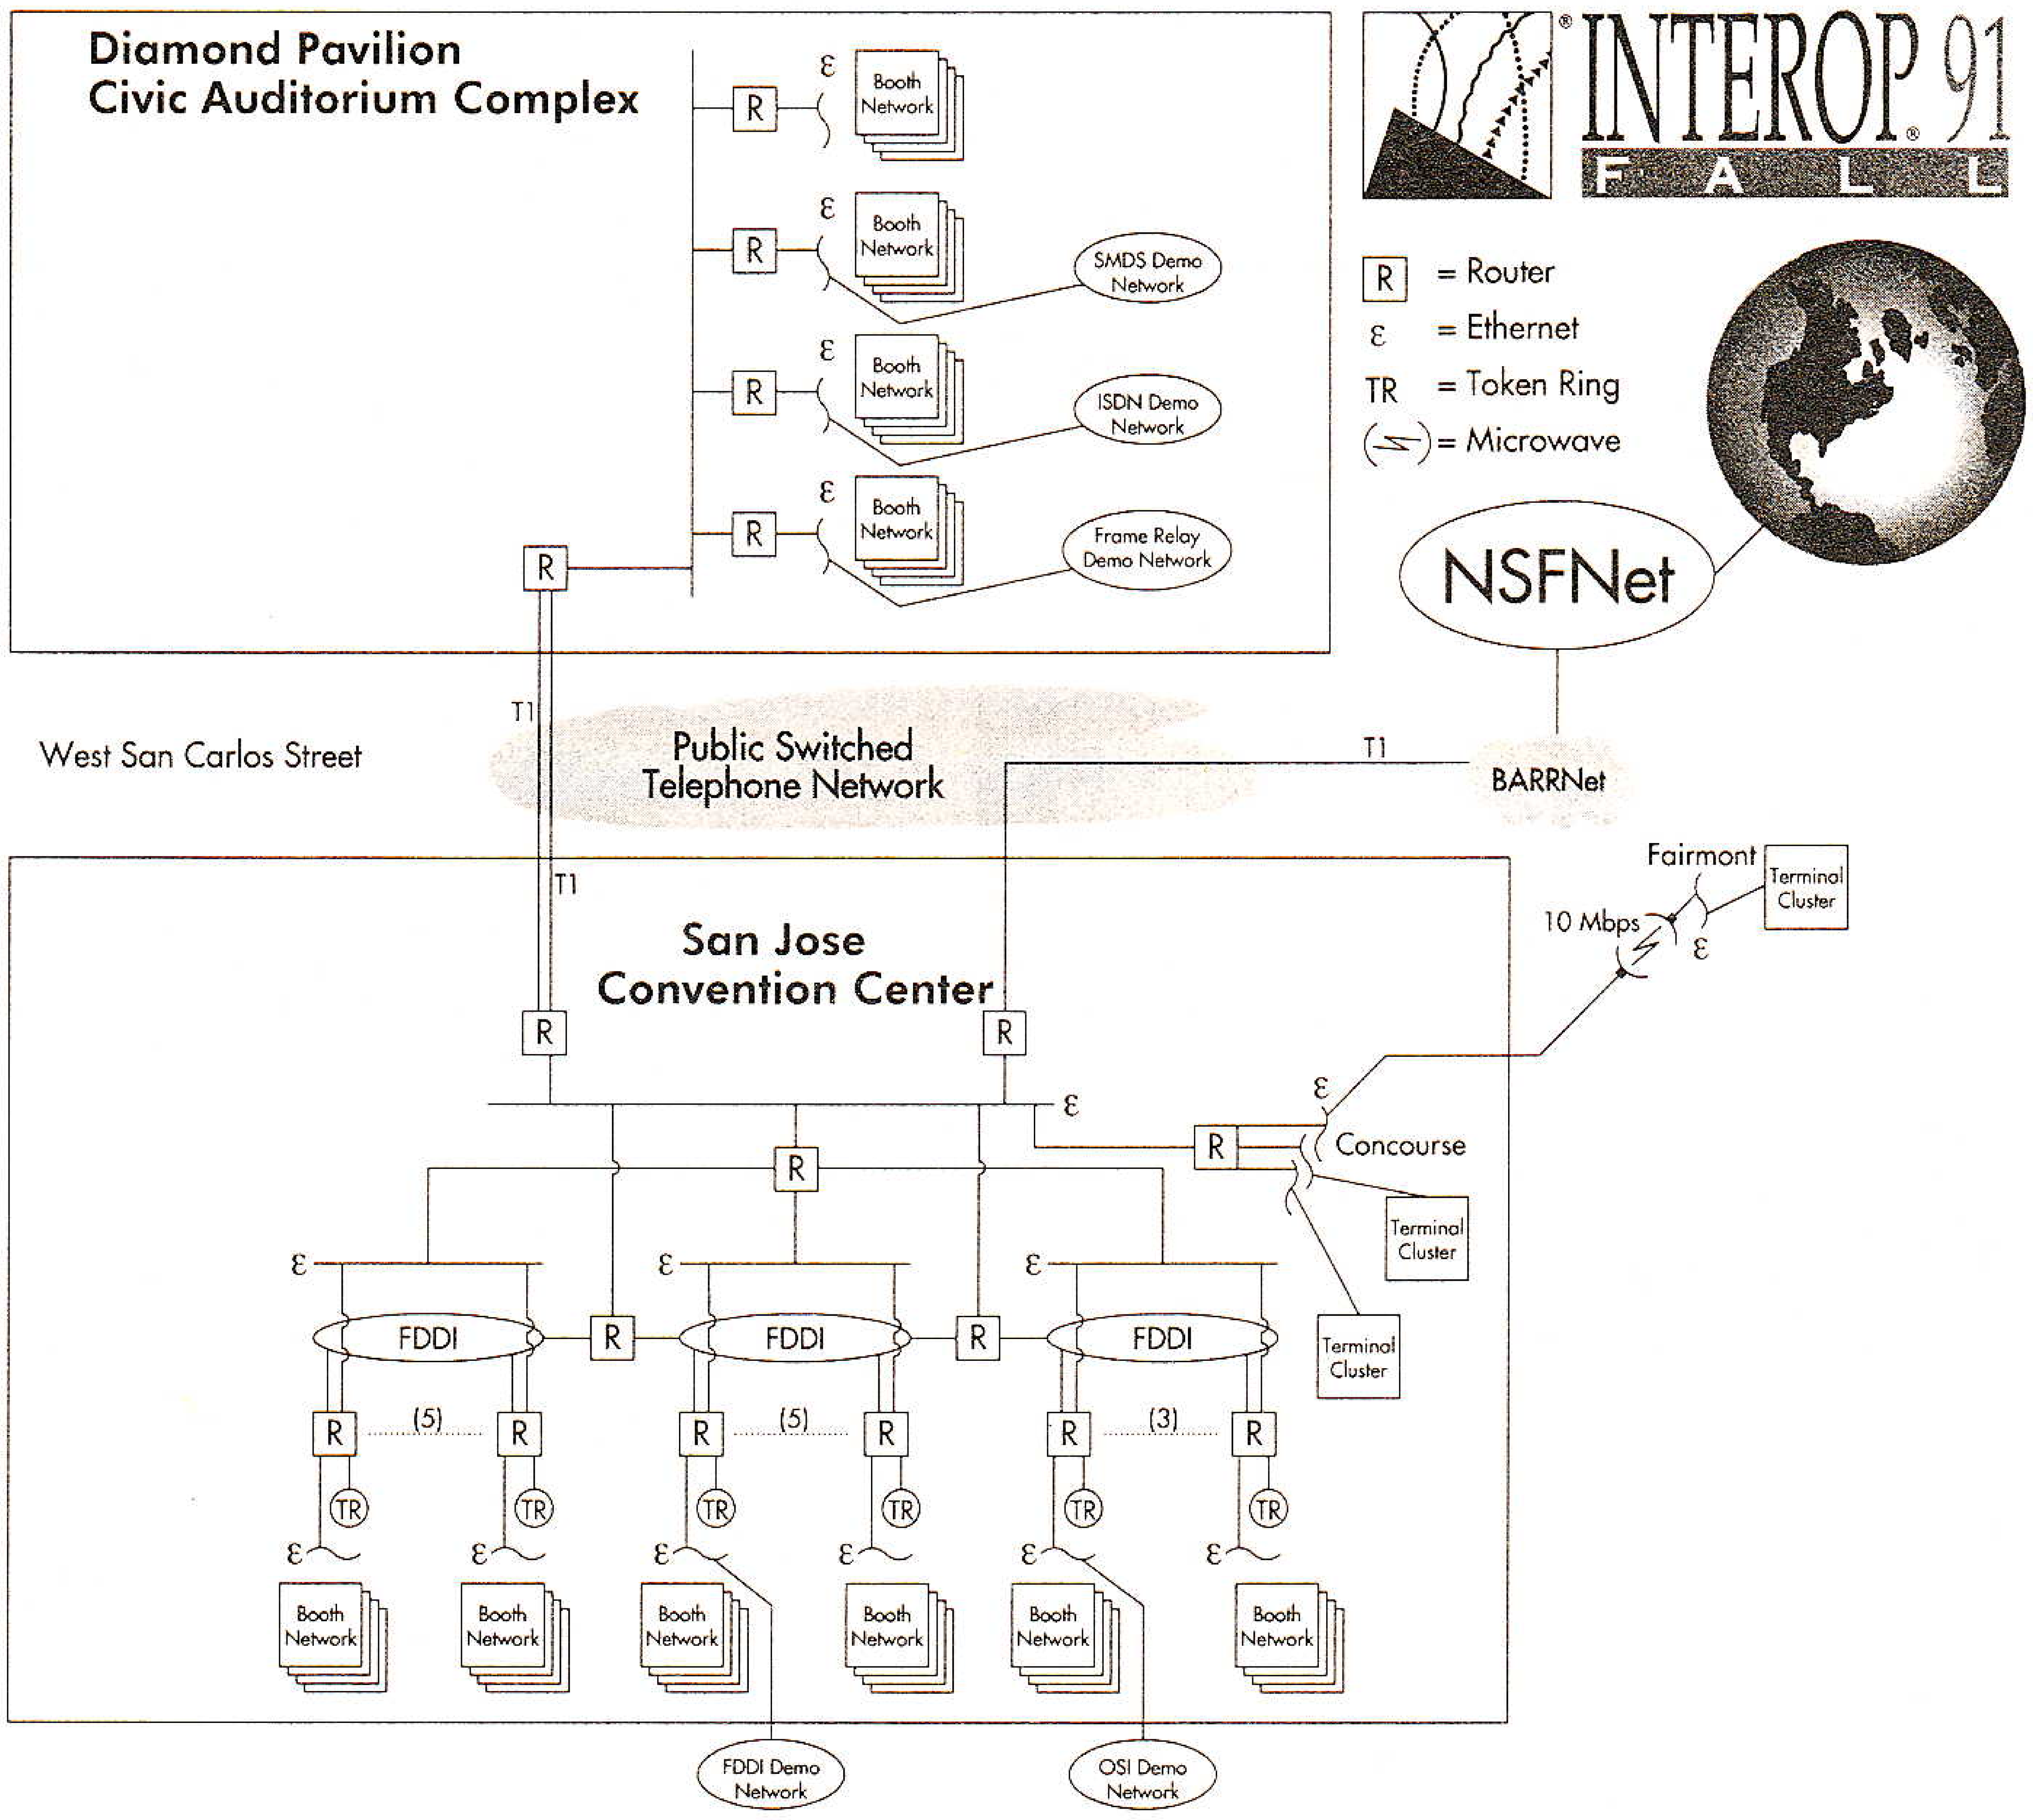
Diamond Pavilion
Civic Auditorium Complex
Booth Network
SMDS Demo Network
ISDN Demo Network
Frame Relay Demo Network
INTEROP 91
FALL
R = Router
ε = Ethernet
TR = Token Ring
= Microwave
NSFNet
West San Carlos Street
T1
Public Switched Telephone Network
BARRNet
Fairmont
Terminal Cluster
10 Mbps
San Jose
Convention Center
Concourse
FDDI
(5)
(3)
FDDI Demo Network
OSI Demo Network

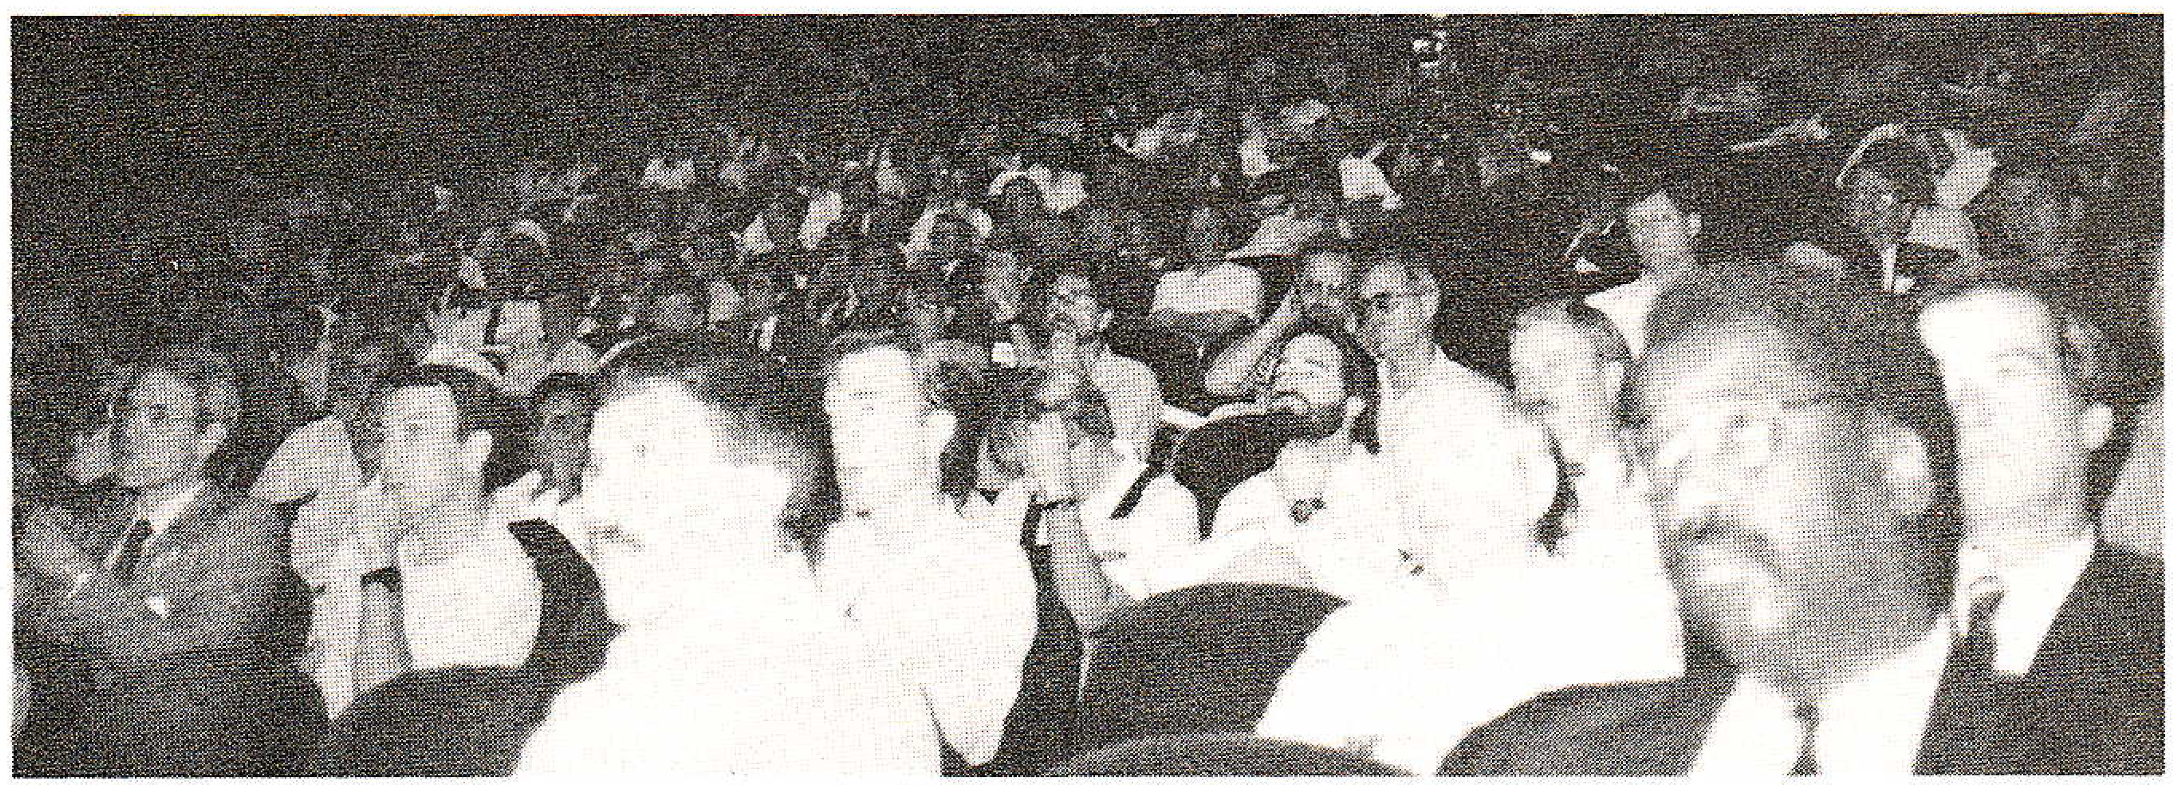


WELCOME TO INTEROP
INTEROP EXHIBITS

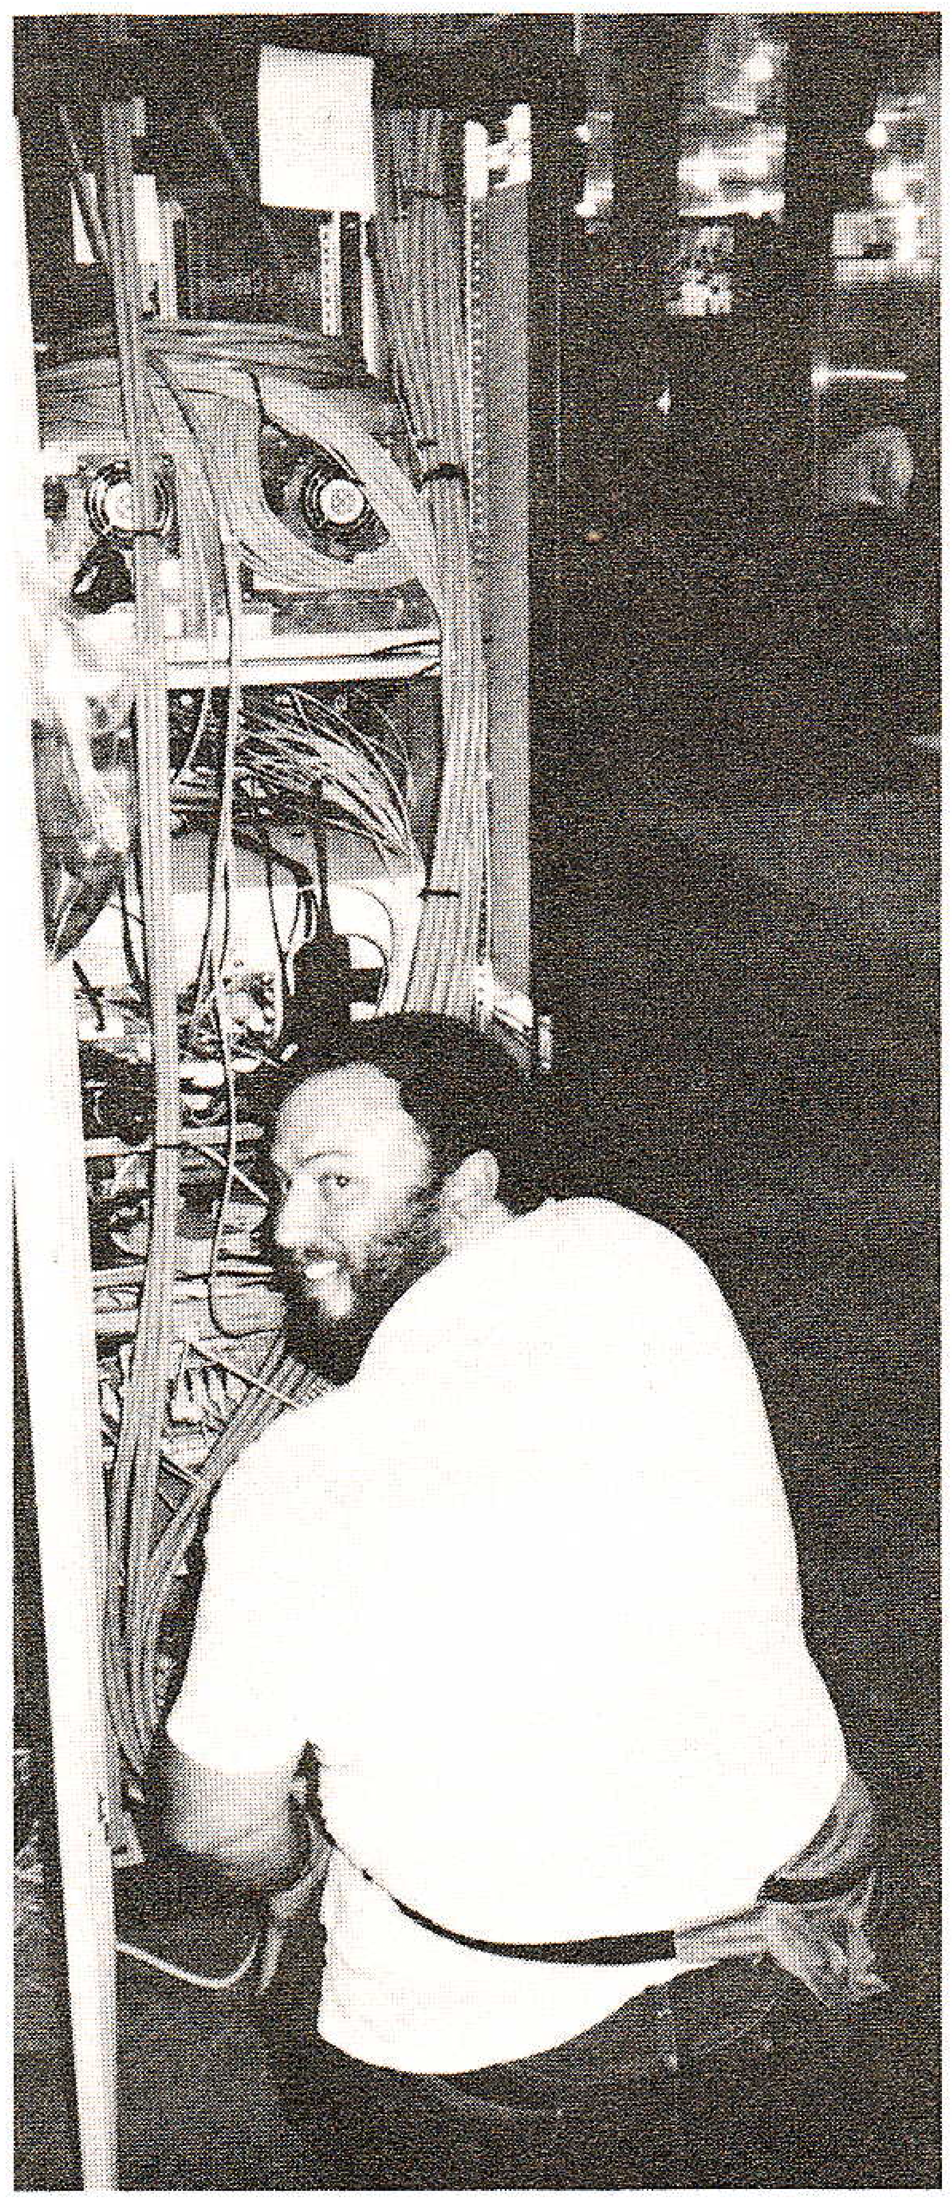

An Overview of OSI NSAP Addressing in the Internet

by Ross Callon, Digital Equipment Corporation

Routing, addressing, and scaling

As the Internet grows, the amount of routing information maintained by routers and passed in routing protocols also grows with it. It is becoming clear that evolution to a ubiquitous worldwide internet, several orders of magnitude larger in scope than the current Internet, is not practical with the current IP addressing scheme.

The OSI Equivalent to IP (the *Connectionless Network Layer Protocol* (CLNP—ISO 8473) makes use of large and flexible hierarchical Network Layer addresses (known as *Network Service Access Point Addresses,* or NSAPs [1]). These addresses provide the flexibility needed to simultaneously solve two critical problems: (i) How to administer a worldwide address space; and (ii) How to assign addresses in a manner which makes routing feasible in a worldwide internet. However, assignment of addresses needs to be done with great care, if the potential advantages of this addressing flexibility are to be realized in actual networks.

The NSAP working group of the *Internet Engineering Task Force* (IETF) has been working on guidelines which describe the principles behind addressing and routing, as applied to assignment of NSAP addresses. Their results have been published as "Guidelines for OSI NSAP Allocation in the Internet" [2]. Although these guidelines are specifically oriented towards the OSI NSAP address space, the same principles are applicable to any network layer address space for use in a large internet.

A lot of research has been done on scaling of routing to very large networks. However, any practical method for use in a large internet requires use of hierarchical network addresses, where the addresses are assigned in a manner which largely corresponds to the network topology.

This requirement makes sense intuitively: In order for routing to scale, at some level there will need to be individual pieces of routing information (for example, single entries in forwarding tables, or pieces of information exchanged in routing protocols) that apply to multiple destinations. This implies that these multiple destinations will in some sense need to be located in "topologically close proximity" in the network.

An example NSAP format

There is a balance that must be sought between the addressing requirements for efficient routing, and the need for decentralized address administration. The NSAP structure from US GOSIP Version 2 is used as an example to illustrate how these two needs might be met. US GOSIP Version 2 defines the following address structure:

	AFI	IDI	DFI	AA	rsvd	RD	Area	ID	Sel
# octets	1	2	1	3	2	2	2	6	1

Here the *Authority and Format Indicator* (AFI) and *Initial Domain Indicator* (IDI) indicate that this is the US GOSIP format. The *DSP Format Identifier* (DFI) indicates that the rest of the address is as illustrated, and may be thought of as an escape mechanism in case a different GOSIP format needs to be defined in the future. The combination of AFI, IDI, and DFI may be thought of as a four octet prefix which indicates that the address is assigned from the US GOSIP address space.

Note: It is not important that the US GOSIP version 2 format be used. For example, the ANSI format under the Data Country Code for the USA and formats assigned to other countries and ISO members may also be used. The GOSIP format is used here only for illustrative purposes.

The *Administrative Authority* (AA) field is assigned by the General Services Administration. It is expected that this field will be assigned to Transit Routing Domains, as well as to organizations which want to receive a "top level" GOSIP address assignment. A *Transit Routing Domain* (TRD) is defined to be a public data network, backbone network, or regional network, which is used to interconnect multiple other routing domains.

However, not all organizations *should* want to receive an AA value. In many cases an AA will be assigned to a TRD, which will then assign RD values to organizations who administer routing domains which are customers of this TRD. If most or all of the customers of a particular TRD use an address based on the AA assigned to that TRD, then the amount of information needed for inter-domain routing can be greatly reduced. In particular, this will allow the set of most or all addresses reachable via the TRD to be abbreviated into a single routing entry.

The *Area, ID,* and *Sel* fields are used to facilitate routing within the routing domain. Generally, a different value for the Area field may be assigned to each area in the routing domain. However, the OSI IS–IS protocol [6, 7, 8] actually uses the combination of all fields up to and including the Area field to indicate the area (and does not care about the exact format of these fields—it just truncates the last "n" octets, where n=7 for domains using US GOSIP addresses, and uses the remaining prefix to indicate the area). This is useful, for example, for routing domains in which addresses from more than one high-level address prefix are used. This, for example, may be useful in the "multi-homed routing domain" case described in [2].

An address is not a route

A network layer address, in some sense, indicates where a particular system is in the network. However, an address does not require any particular route be used to reach that system. It is useful to consider how the address used to identify a particular system may effect the method used to route to that system.

Let's suppose that there is a US-based corporation (which we will call "X"), which runs its own internal network. X is attached to two NSFNET regional networks (which we will call "A" and "B"). Let's suppose that X has several options for how to get an address:

1) X could apply directly to OSI, and get a top-level address. The address for corporation X would essentially say: "organization-X."

2) X could apply to GSA, and get an AA identifier from the US-GOSIP space. Thus the address for X would say: "US-GOSIP/organization-X."

3) X could similarly apply to ANSI, and get an address which says: "US-ANSI/organization-X."

4) X could apply to one of the regionals to which it is attached (let's say regional A) and get an address from the prefix assigned to that regional. Let's assume that regional A previously got an AA assignment from the US GOSIP space, and that A is assigning RD values to its customers from this address space. Thus the address for X would say: "US-GOSIP/regional-A/organization-X."

continued on next page

OSI NSAP Addressing in the Internet *(continued)*

In each of these cases, there is a single prefix which describes all addresses reachable within organization X. Thus it is in principle possible for all TRDs worldwide to maintain a routing entry for organization X. Provided that all TRDs maintain an explicit route to organization X, it makes no difference to routing which of these solutions is chosen. Thus, none of these solutions preclude maintaining an optimal route.

The different effects on routing apply only if some of the TRDs are not willing to maintain a route to organization X. Given that a worldwide internet could potentially have hundreds of thousands, or even millions of routing domains, it is unlikely that any single TRD will in fact maintain individual routes to all routing domains worldwide.

With solution 1, the prefix for X is top-level, so that if a TRD is not willing to maintain a route to X, then packets sent via that network to a destination in X are not deliverable.

With solutions 2, 3, and 4, the prefix for X is US-based. Thus, any TRD outside of the US, in case of doubt (i.e., if it did not maintain a special entry for organization X), could take a guess and forward the packet to the US. This will allow delivery of the packets if the networks in the US all know about organization X, and if they are interconnected.

Again, if the TRDs in the US all maintain an explicit route entry to X, then once a packet makes it to the US, solutions 2, 3 and 4 work the same. The difference only comes about if not all TRDs in the US do this. In particular, with solutions 2 and 3, the prefix for X just says "this address is in the US, and is organization X." Thus, any TRD in the US that does not have an explicit entry for X is lost. With solution 4, the address basically says "this address is in the US, has an attachment to regional A, and is organization X." Thus if you have a route entry for this organization you take the best route there (which might be via regional A, might be via regional B, or might be via some other route). If you don't have an entry for organization X, but you do have an entry for regional A, then the packet will be routed via regional A.

The address chosen for a particular destination therefore does not require that any particular route is chosen. However, if the address is assigned in a manner which is related to the topology of the internet, then the address does provide a sort of "hint," which can be used to route packets in the absence of better information.

Where do I get an address?

So, where does a routing domain administrator go to get an NSAP address?

For routing domains which are not attached to any transit routing domain, there are a lot of places that you can go to get a unique valid NSAP address. You can obtain an address assignment from national standards bodies, or can create a globally unique NSAP address from a telephone number or telex number assigned to you (see below).

For routing domains which are attached to a single transit routing domain the best bet is to use an address taken from the space administered by that transit routing domain. This minimizes the effort required (since the folks from the transit routing domain should do most of the work of getting you a valid globally-significant NSAP address prefix), and this also makes inter-domain routing more efficient.

In addition, as described above, this increases the possibility that packets destined to your domain will be delivered successfully.

For routing domains which are attached to multiple transit routing domains, there is no single "best" choice for an NSAP format. Such routing domains may use a "top level" address prefix, may use a prefix chosen according to one of their attachments, or may use several prefixes (one for each transit routing domain to which they are attached). A set of possible choices, as well as advantages and disadvantages of each choice, is described in more detail in [2].

The OSI NSAP authorities and formats

The NSAP address standard allows addresses to be assigned according to a number of different authorities. This section outlines each one briefly, and describes some possible situations under which each is likely to be most useful.

- *Local:* The ISO NSAP standard allows for a local NSAP space. NSAPs allocated from this space are not globally unique. These addresses are reasonable for isolated routing domains. However, if your network is to be interconnected with other OSI networks at any time, you should not use the private network format.

- *ISO DCC:* This address format makes use of the CCITT *Data Country Code* (DCC) to assign a unique address prefix to each country. A national standards body for each country then assigns a prefix value from the national space to organizations within that country.

This space may be used in two cases: (i) Addresses under this space may be assigned to transit routing domains within a country, allowing organizations which are customers of those transit routing domains to obtain addresses under that transit routing domain's prefix; (ii) Organizations may obtain address assignments directly from their national standards organization.

- *ISO 6523-ICD:* ISO 6523 provides an international standard for identification of organizations. This includes an *International Code Designator* (ICD) which is assigned to ISO members and liaison organizations, as well as some organizations of international importance. For example, an address prefix from this address space has been assigned to the US *National Institute of Standards and Technology* (NIST), and is administered by the General Services Administration for GOSIP Version 2 addresses.

- *X.121:* This address format is based on the X.121 addresses used with X.25 networks. This address format may be used with any valid X.121 address that has been assigned to your organization.

The node whose X.121 address you use does not actually have to be part of the network you are creating an NSAP address prefix for, because it is serving only as a unique number that has been assigned to you. However, if you use this type of address, other organizations may use the X.121 address as a "hint" when routing traffic to nodes in your network. This method therefore works best when the X.121 address that you use is the address of your network's connection to the public data network.

- *F.69:* This format makes use of an F.69 telex number to generate unique NSAP address prefixes. This allows unique NSAPs to be generated based on any valid telex number, but is not likely to provide any useful hints for routing as described above.

continued on next page

OSI NSAP Addressing in the Internet *(continued)*

- *E.163:* This is again similar to the F.69 format, except that an E.163 international telephone number is used. Again, this is useful primarily to allow a unique NSAP address prefix to be generated.

- *E.164:* This allows an NSAP to be based on an E.164 Integrated Services Digital Network (ISDN) number. This is similar to the X.121 address format. If you use this type of address, other organizations may use the E.164 address as a "hint" when routing traffic to nodes in your network. This method therefore works best when the E.164 address that you use is the address of your network's connection to a public ISDN network.

More information

For more information, the reader is directed to the references. "Guidelines for OSI NSAP Allocation in the Internet" [2] is recommended as an overview of the routing and scaling issues. "American National Standard for the Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address" [4] and "Government Open Systems Interconnection profile (GOSIP) Version 2" [5] specify the format to be used with ANSI-assigned and US GSA/GOSIP addresses, respectively.

References

[1] "Network Service Definition, Addendum 2: Network Layer Addressing," International Standard 8348/Addendum 2, 1988.

[2] "Guidelines for OSI NSAP Allocation in the Internet," R. Colella, E. Gardner, and R. Callon, RFC 1237, July 1991.

[3] "The OSI Network Layer Addressing Scheme, Its Implications, and Considerations for Implementation," Christine Hemrick, NTIA Report 85-186, US Department of Commerce, National Telecommunications and Information Administration, 1985.

[4] "American National Standard for the Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address," (pending final approval by ANSI).

[5] "Government Open Systems Interconnection Profile (GOSIP) Version 2," GOSIP Advanced Requirements Group, Federal Information Processing Standard 146–1, U.S. Department of Commerce, National Institute of Standards and Technology, April 1991.

[6] "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473)," ISO DIS 10589, November 1990.

[7] Tsuchiya, P., "Component of OSI: IS–IS Intra-Domain Routing," *ConneXions,* Volume 3, No. 8, August 1989.

[8] Callon, R. & Perlman, R., "The Great IGP Debate—Part One: IS–IS and Integrated Routing," *ConneXions,* Volume 5, No. 10, October 1991.

ROSS W. CALLON is with the Distributed Systems Architecture group at Digital Equipment Corporation in Littleton, Massachusetts. He is working on OSI–TCP/IP Interoperability issues, and is the architect for the Integrated IS–IS protocol. Mr. Callon is also the co-area director of the OSI Integration area of the IETF, and is the chair of the IETF IS–IS working group. Mr. Callon received his B.Sc in Mathematics from the Massachusetts Institute of Technology, and his M.Sc in Operations Research from Stanford University.

The ITU Adopts a New Meta-Standard: Open Access

by Carl Malamud

Major announcement

INTEROP 91 Fall featured a major announcement by Dr. Pekka Tarjanne, Secretary-General of the International Telecommunication Union. Using a live video teleconferencing link from Geneva, Dr. Tarjanne told INTEROP conference attendees that he has decided to allow the Internet community to post all ITU standards for distribution, at no charge, over the network.

The ITU standards are a crucial body of international standards, ranging from X.25 to G3/G4 fax to Signalling System 7 to X.400 messaging and the V series modem definitions. The CCITT standards set, known as the *Blue Book*, is over 19,000 pages long.

File server

The standards are initially available on a server donated by Sun Microsystems and maintained by the University of Colorado. Many other sites including UUNET will house copies of the standards archive.

To obtain standards from the server, users can initiate an anonymous FTP session to `digital.resource.org` (the preferred address) or `bruno.cs.colorado.edu` (the alias). Electronic mail sent to `infoserve@digital.resource.org` with the world "help" in the message body will return instructions on how to use a sophisticated mail-based archive server. Comments on any aspects of this program may be sent to `standards@digital.resource.org`.

Conversion

The ITU maintained the Blue Book on a Siemens mainframe using a 1970s style proprietary text formatting system (complete with their own character set named *Zentec*). A conversion program was written in *perl* which is able, with some notable exceptions, to convert the data into more rational formats.

The conversion program, in its initial implementation, converts the Blue Book into *troff.* For convenience, ASCII (i.e., *nroff*) and *PostScript* (i.e., *psroffed*) versions of the standards are posted along with their *troff* source.

Limitations

Two notable problems will be apparent in the converted documents. First, some tables and formulas were not able to be converted due to somewhat incomplete documentation on the original format (and the limited time and skills of the conversion programmer).

The second major limitation is the lack of integrated graphics. The ITU maintained graphics in Autocad, but manually added all text with typewriters, glue, and similar anachronisms. As an initial workaround, close to one thousand figures were manually scanned and are posted in TIFF and EPS formats. It is hoped, in the future, that we can provide a more elegant solution.

In addition to the Blue Book, there are quite a few more recent standards stored in other formats. These formats include Microsoft Word for Windows (the new ITU publishing platform), Rich Text Format (a Microsoft revisable form standard), ASCII, Word Perfect, and a few Samna files (the ITU precursor to Word).

Although there are certainly deficiencies in the free versions of the ITU standards, it is hoped that the Internet community will look beyond the individual bytes to the symbolic nature of this important announcement.

continued on next page

ITU Adopts a New Meta-Standard *(continued)*

Cost

The ITU and ISO have, until now, only made their standards available by paper at great cost. The ITU derived annual revenues of SFr 8 million (roughly US$ 5 million) from document sales. ISO also makes a considerable amount from document sales although they refuse to divulge specific revenue or cost figures, claiming the information is "proprietary."

High cost has meant that the most important communities—individuals who will implement the standards—have not had ready access to these vital documents.

Copyright

Many considered one of the most surprising aspects of the INTEROP session to be a speech, made via the video link from Geneva, by Anthony Rutkowski, one of the senior lawyers at the ITU. Rutkowski presented a detailed analysis of the legal basis for asserting copyright on standards documents. He concluded that it was highly doubtful if international organizations (such as ISO or the ITU) would be able to successfully assert copyright protection in a court of law over the content of standards. (The video link was provided courtesy of US Sprint and Compression Labs Incorporated).

A first step

The announcement by the ITU is a radical change in policy and represents the new leadership of Dr. Pekka Tarjanne. Dr. Vinton Cerf, the chairman of the IAB, underscored the significance of these new policies when he informed the INTEROP audience that he had received calls from the White House and the FCC wanting to know more details.

The ITU announcement is an important step, but it is only the first step. ISO, ANSI, the IEEE, Bellcore, and all other standards-making and standards-coordination bodies must firmly endorse the principle that the results of the standards process should be easily accessible at low cost or no cost.

In addition, bodies like ISO and ANSI should investigate their processes to see if there are means that can, while preserving the vital principle of due process, enhance the speed and relevance of standards making. An important first step would be to post all working documents on the Internet.

Posting documents on the Internet is technically feasible. In less than one month, a very small group of volunteers were able to convert most of the ITU standards set (not to mention scanning in images, setting up the hardware and installing support software). There are no technical reasons to not post standards—all objections are based on a political desire to retain control or a financial desire to enhance revenue.

Objection

The most common objection by groups like ANSI to posting standards is that document sales "fund the process." While ANSI and ISO refuse to divulge their cost and revenue structure, it is certainly true that under the current procedures the document sales are important to both groups.

However, selling standards at very high prices undermines the very purpose for which groups like ISO and ANSI were formed. It is as if the American Cancer Society were to sell cigarettes as a way of funding their work. ISO and ANSI were formed to promote the widespread acceptance of standards: their current policies severely undermine those goals.

While the revenues from document sales are important, it should be noted that these sums are a mere pittance when compared to the tremendous donation in time and money by the voluntary participants in the standards process. Making standards available at very high prices undermines the tremendous donation by vendors and users to making standards.

Conclusion

Funding any non-profit activity is always difficult. If ANSI and ISO open up, the talented people who help make the standards can come up with creative solutions to financing that will maintain wide distribution of standards and not force ANSI to hold periodic bake sales. First, however, standards bodies must adopt the meta-standard of *availability*.

References

[1] Malamud, C., "Letter to the Editor," *ConneXions,* Volume 5, No. 2, February 1991 (p 20–21).

[2] Rutkowski, A. M., "Networking the Telecom Standards Bodies," *ConneXions,* Volume 5, No. 9, September 1991.

[3] Chappell, D. "Components of OSI: A taxonomy of the players," *ConneXions,* Volume 3, No. 12, December 1989.

[4] Rose, M., "The Open Book: A Practical Perspective on OSI," Prentice-Hall, ISBN 0-13-643016-3, 1990.

[5] Rose, M., "The Simple Book: An Introduction to Management of TCP/IP-based internets," Prentice-Hall, ISBN 0-13-812611-9, 1990.

[6] Rose, M., "The Little Black Book: Mail Bonding with OSI Directory Services," Prentice-Hall, ISBN 0-13-683210-5, 1990.

CARL MALAMUD is one of the founders of the *Document Liberation Front,* a volunteer group dedicated to seeing standards available for free on the Internet. ANSI and ISO officials can reach him as `carl@malamud.com`.

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Want to enquire about back issues? (there are now more than fifty-five to choose from; ask for our free 1987–1991 index sheets). We want to hear from you. Simply write, call, fax, or e-mail to:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040–1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: `connexions@interop.com`

Call for Participation: INET '92

The first international conference organized by the *Internet Society* (just being formed) focusing on worldwide issues of research and academic networking—*INET '92*—will be held on June 15–18, 1992 in Kobe, Japan. The goal of this conference is to bring together individuals from university, industry and government who are involved with planning, developing, implementing, managing and funding national, regional and international research and academic networks.

Sponsors

The conference is sponsored by The Internet Society in cooperation with ACM SIGCOMM, IPSJ (Information Processing Society of Japan), CREN, EDUCOM, FARNET, IAB and NetNorth in North America, JAIN, JCRN, TISN, and WIDE in Japan, EARN, EUUG (EUnet), NORDUNET and RARE in Europe.

Topics

The conference agenda will include plans and status reports on research and academic networks throughout the world. In addition, possible topics for conference sessions include, but are not limited to, the following:

Technology and Services:

- Progress toward international open network protocols
- Interoperability among existing national & international networks
- Collaboration technologies
- Security, management and authentication in managing networks
- Technologies of the 90s
- Mail and directory services
- Very high speed networks
- Data storage: terabytes and beyond
- Networks and network services of the 21st century

Applications:

- Network support of international scientific collaboration
- Network access to scientific papers and data across national boundaries
- Supercomputing
- High energy physics
- Workstation teleconferencing
- Computer supported collaborated works
- Education
- Libraries

Regional Issues:

- Asia-Pacific Rim
- Eastern Europe
- Europe
- Latin America
- North America
- Africa
- Special Issues for the Third World

Policy Issues:

- Globalization of services
- Commercialization/privatization
- Coordination of international links
- Appropriate use
- Security
- Privacy
- Telecommunications policy

Paper submissions

For Technology & Services and Applications sessions, please submit 6 copies (in English) of double-space typed manuscript (maximum of 20 pages) with an abstract to:

Prof. Haruhisa Ishida
Computer Centre, University of Tokyo
Bunkyo-ku, Tokyo
Japan 113
Phone: +81-3-3818-0287
FAX: +81-3-3814-7279

You may also submit a *PostScript* version of your paper by e-mail to `ishida@u-tokyo.ac.jp`.

Tutorials

In addition to papers, proposals for one-day tutorials are solicited in any of the conference areas. Such proposals should be submitted to the program chair or Deborah Estrin (`estrin@usc.com`) by January 6, 1992.

Important dates

January 6, 1992:	Manuscript due and proposal due for a tutorial
March 1, 1992:	Notification of acceptance to authors
April 15, 1992:	Camera-ready papers due

Workshop

A one-day workshop for representatives of developing and third world countries, "How to Plan for and Implement Research and Academic Networks," will be held in conjunction with INET '92.

For more information

To be added to the conference mailing list, send mail, FAX or e-mail to one of these addresses:

Japan & Pacific Rim:

WIDE project		
INET '92 secretariate	Phone:	+81 466 48 9433
5322 Endo, Fujisawa	FAX:	+81 354 90 7002
Kanagawa 252, Japan	E-mail:	`inet92@wide.ad.jp`

North America:

EDUCOM		
INET '92 secretariate	Phone:	+1 202 872 4200
1112 16th Street, NW	FAX:	+1 202 872 4318
Washington DC 20036, USA	E-mail:	`inet92@educom.edu`

Europe:

UNI-C		
INET '92 secretariate	Phone:	+45 45 93 83 55
DTH, bygning 305	FAX:	+45 45 93 02 20
DK-2800 Lyngby, Denmark	E-mail:	`inet92@vm.uni-c.dk`

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED